Tardos fingerprinting is better than we thought

B. Škorić, T.U. Vladimirova, M. Celik, J.C. Talstra

Abstract

Tardos has proposed a randomized fingerprinting code that is provably secure against collusion attacks. We revisit his scheme and show that it has significantly better performance than suggested in the original paper. First, we introduce variables in place of Tardos' hard-coded constants and we allow for an independent choice of the desired false positive and false negative error rates. Following through Tardos' proofs with these modifications, we show that the code length can be reduced by more than a factor of two in typical content distribution applications where high false negative rates can be tolerated. Second, we study the statistical properties of the code. Under some reasonable assumptions, the accusation sums can be regarded as Gaussian-distributed stochastic variables. In this approximation, the desired error rates are achieved by a code length twice shorter than in the first approach. Overall, typical false positive and false negative error rates may be achieved with a code length approximately 5 times shorter than in the original construction.

Keywords: watermark, fingerprinting, collusion, coalition, traitor tracing, random code.

1   Introduction

1.1   Digital fingerprinting

Digital content, such as songs, photographs or movies, can be copied inexpensively without any loss of quality. Moreover, these copies can be easily redistributed without permission from the original rights holders. According to [RW2004], unauthorized sharing of music on peer-to-peer (P2P) networks in college campuses reduces the potential revenue of the recoding industry by as much as 20%.

One way of countering unauthorized redistribution is to uniquely mark each individual instance of the originally distributed content, so that the recipient ('user') can be identified if that content appears on a P2P network. The authorized distributor (or the content owner) embeds a unique mark, also called a 'forensic watermark' or a 'fingerprint', into each instance of the content before transmitting it to the user. The embedding algorithm ensures that the mark is imperceptible, i.e. the quality of the content is not degraded by the mark. Moreover, the location and nature of the mark is kept secret from the user to prevent him from locating and altering the mark. Typically the mark consists of a set of symbols from a q-ary alphabet, where each symbol is embedded into a different part of the content, e.g. different scenes in a movie. When an unauthorized copy is found, the content owner, knowing all the details, can detect the mark and identify the source of the unauthorized copy. Forensic watermarking has already been successfully applied in practice [CNN2004].

1.2   Collusion resistance

A group of recipients (called 'colluders' or 'a coalition') can collaborate to escape identification. Comparing their content copies, they can find the locations where their content, and thus their marks, differ. These locations are called the 'detectable positions'. By cleverly manipulating the

content at those locations, the colluders can attempt to create a version of the content that cannot be traced back to any of them. Such an attack is called a *collusion attack*.

Collusion attacks and fingerprinting schemes that show resistance to these attacks have been studied since the late 1990's. The often used *marking condition* assumes that the colluders are unable to change the symbols (marks) in undetectable positions. Under the marking condition, one can distinguish between the following attack models, which differ in the type of manipulation the attackers are allowed to perform:

- The *restricted digit model* or *narrow-case model* allows the colluders only to 'mix and match', i.e. to replace a symbol in a detectable position by any of the symbols they have received in that position.

- The *unreadable digit model* allows for slightly stronger attacks. The attackers can also introduce an unreadable symbol '?' in detectable positions.

- The *arbitrary digit model* allows for even stronger attacks. The attackers can put any (arbitrary) q-ary symbol (but not the unreadable symbol '?') in the detectable positions.

- The *general digit model* allows the attackers to put any symbol, including the unreadable symbol '?', in the detectable positions.

In the case of a binary alphabet all four attack models are equivalent. The content owner can map the unreadable symbol '?' to either of the binary symbols without loss of generality.

Video fingerprinting applications face a number of severe constraints in practice. First of all, there is a limit on the number of locations (m) suitable for watermark embedding. A typical fingerprinting system can reliably extract approximately seven bits per minute of video content [DCI2007]. Furthermore, constraints on decoding complexity and perceptual quality limit the number of different symbols that can be embedded in each location. Hence, the alphabet size q for a fingerprint code is limited (typically $q \leq 16$). Finally, mass market content distribution systems need to accommodate a very large number of users (e.g. millions or even hundreds of millions). Under these constraints, the authorized distributor is interested in the fingerprint code which can resist the largest coalition size (c_0).

In the last decade, various fingerprint codes have been proposed. Some of these codes are deterministic, i.e. they can identify at least one member of the coalition with certainty, without the danger of accusing an innocent user. For instance, Identifiable Parent Property (IPP) codes proposed in [HvLLT1998] are deterministic. However, the scheme is limited to a coalition size of two. In [SSW2001], Staddon et al. proved the existence of a deterministic fingerprinting code which is resistant against c_0 colluders. The code is of length $m = c_0^2 \log_q(n)$, where n is the number of users. However, it requires an impractically large alphabet size, $q \geq m - 1$. Another deterministic scheme, presented in [CFNB2000], has a similar length $m = 4c_0^2 \log n$ with a smaller alphabet size $q = 2c_0^2$. Still, the alphabet size quickly becomes prohibitive for mass market content distribution systems.

When the application can tolerate a nonzero probability of error, randomized fingerprinting codes with smaller, even binary, alphabets can be used. In a typical fingerprinting application, the most important type of error is the False Positive (FP) error, where an innocent user gets accused. The probability of such an event must be extremely small; otherwise all accusations become dubious, making the whole fingerprinting scheme unworkable. We will denote by ε_1 the probability that a specific innocent user gets accused. The notation η is used for the probability that there exist innocent users among the accused. The second type of error is the False Negative (FN) error, where the scheme fails to accuse any of the colluders. In practical content distribution applications, fairly large FN error probabilities can be tolerated, as the content owner can collect evidence from multiple pieces of content over a period of time. We denote the FN error probability by ε_2.

In [BS1998], Boneh and Shaw presented a binary ($q = 2$) randomized code with length $m = \mathcal{O}(c_0^4 \ln \frac{n}{\eta} \ln \frac{1}{\eta})$, which uses concatenation of a partly randomized inner code with an outer code. They also proved a lower bound on the required length for any binary code that is resistant against

c_0 colluders: $m \geq \frac{1}{2}(c_0 - 3) \ln \frac{1}{c_0 \eta}$. In [PSS2003], Peikert et al. proved a tighter lower bound for a restricted class of codes with a limited number of 'column types': $m = \Omega(c_0^2 \ln \frac{1}{c_0 \eta})$.

In [Tar2003], Tardos further tightened the lower bound for the arbitrary digit model and the unreadable digit model: $m = \Omega(c_0^2 \ln \frac{1}{\varepsilon_1})$ for arbitrary alphabets. In the same paper, he described a fully randomized binary fingerprinting code achieving this lower bound. This code has length $m = 100 c_0^2 \lceil \ln \frac{1}{\varepsilon_1} \rceil$.

While [Tar2003] proposed a practical code with optimal behavior for large c_0, it leaves a number of open questions:

- The author uses a number of arbitrary-looking constants in his proofs, such as 20, $\frac{1}{10c_0}$ $\frac{1}{20c_0}$, $\frac{1}{300c_0}$. Similarly, the constant 100 appears in the minimum code length expression $m = 100 c_0^2 \lceil \ln \frac{1}{\varepsilon_1} \rceil$. While these numbers allow for important properties to be proven, it is not at all clear if they have been chosen in an optimal way.

- He also makes seemingly arbitrary choices for the accusation weight function—which specifies how strongly to accuse a user per symbol if his symbol is equal to the symbol found in a pirated copy—and the distribution function—which specifies probabilities used in generating the random code words,—hinting that they are optimal, but not providing a proof.

- In the proofs, the FN error probability ε_2 is coupled to the FP probability ε_1. While Tardos remarks that they can be decoupled, it is not clear how each exactly influences the code length on its own. Furthermore, the coupling is such that FN rate is much smaller than the FP rate, $\varepsilon_2 \ll \varepsilon_1$. As mentioned earlier, in practical applications the opposite $\varepsilon_2 \gg \varepsilon_1$ may be desirable. This opposite case is not studied.

1.3   Contributions and outline

In this paper, we provide answers to the aforementioned issues with Tardos' construction, which were left open in [Tar2003].

- In Section 2.2, we generalize the Tardos' construction, introducing variables in place of numerical constants and generic functions instead of the functions specified in [Tar2003]. We state 'Soundness' and 'Completeness' properties, which specify the desired FP and FN error conditions. Our approach is similar in spirit to Hagiwara et al.'s in [HHI2006], but our results are not restricted to small coalitions.

- In Section 3.1, we state the conditions on the construction parameters for a scheme that satisfies both the 'Soundness' and 'Completeness' properties. These conditions are derived in Sections 3.2 and 3.3, respectively. We employ a proof method very similar to [Tar2003], but we specifically do not couple FP and FN error rates. In Section 3.4, the results of the preceding subsections are combined to arrive at a condition for the code length.

- In Section 3.5, we show that Tardos' choice for the accusation weight function is optimal. In Section 3.6, we show that his choice for the distribution function used in generating the random code words is optimal within a limited class of functions.

- In Section 3.7, we arrive at the smallest possible code length parameter value, specific for our proof method, that allows for Soundness and Completeness. In the case of large c_0, this value lies slightly above $4\pi^2$, which is a significant improvement over the original value '100'.

- In Section 4, we present the results of a numerical search for model parameters that yield the shortest possible code length. For certain realistic choices of n, c_0, ε_1 and ε_2, we find that the code length can be reduced by a factor of two or more. For large ε_2 our theoretical large-c_0 result seems to be approached.

- In Section 4, we further see that the code length depends only weakly on ε_2, as was also mentioned by Tardos. Nonetheless, we observe a significant advantage for decoupling the FP and FN error probabilities. When Tardos' coupling of $\varepsilon_2 = \varepsilon_1^{c_0/4}$ is enforced, the code length constant '100' can be reduced only to approximately 90. When the rates are decoupled and the FN rate is allowed to increase, for fixed ε_1 and c_0 values, the constant can be reduced to approximately 45.

We also study the statistical properties of the scheme. The results give us an insight into the 'average' behavior of the scheme and are much simpler than the expressions involved in the formal proofs.

- In Section 5.1, by modeling the accusation sums as normally distributed stochastic variables (an approximation motivated by the Central Limit Theorem), we state a simple approximate condition for the code length based on the coalition size, FP and FN error rates.

- In Sections 5.2 and 5.3, we compute the mean and the variance of the accusation sums for innocent users and the coalition, without any assumptions on their distribution. Similarly, we derive conditions on the code length and the accusation threshold based on the desired false positive and false negative error rates, in Section 5.4.

- In Section 5.5, we identify an 'extremal' colluder strategy. It maximizes our expression for the minimally required code length. The strategy is to output a '1' whenever this is allowed by the marking condition.

- In Section 5.6, we assume that the probability distributions of the accusation sums are perfectly Gaussian. Using this approximation, we show that (for large c_0) error rates ε_1, ε_2 can be achieved with code length $m \approx 2\pi^2 c_0^2 \ln \frac{1}{\varepsilon_1}$. The Gaussian approximation is justified in Appendix C.

2   Tardos revisited

2.1   Generalization

In [Tar2003], Tardos proposed a randomized fingerprinting code resilient against collusion attacks. His fingerprinting scheme is particularly known for its short code length. Nonetheless, due to a number of implicit parameter choices in [Tar2003], it is not clear if his explicit construction achieves the shortest possible code length allowed by the scheme. In this section, we generalize Tardos' construction in anticipation of our study in the following section. In particular, we make three generalizations. First, we replace various fixed numerical parameters by variables and investigate the conditions on these variables that allows us to still carry out the security proofs. This helps us to modify the system parameters so as to obtain even shorter code lengths. Second, we allow for the desired false positive and false negative error probabilities to be chosen independently. As opposed to the original construction where these probabilities were coupled, this generalization allows us to better align the scheme to practical content fingerprinting requirements. Third, we do not assume any specific form for the functions f, g_1 and g_0 (see Section 2.2) in Tardos' construction. The introduction of arbitrary functions f, g_1, g_0 leads to a lot of extra effort in carrying out the proofs. However, with this generalization, we can show that Tardos' choices for g_1, g_0 are optimal (for the proof technique employed in [Tar2003]), and that his choice for f is likewise optimal within a specific class of smooth functions.

2.2   Notation and definitions

We describe our generalized version of the binary Tardos fingerprinting scheme below. We adhere to the notation in [Tar2003] whenever possible. Moreover, we denote the explicit parameter choices made by Tardos by the superscript 'T' to avoid confusion.

The fingerprinting scheme has n recipients ('users'). Each user is assigned a codeword of length m. The set of colluders (the coalition) is denoted as C, whereas the number of colluders is denoted as c. The coalition size up to which the scheme has to be resistant is denoted as c_0. The content owner generates an $n \times m$ matrix X; the j-th row of X is the codeword embedded in the content of user j. The part of X received by the coalition is denoted as X_C. The colluders use a 'C-strategy' ρ to produce an m-bit string $y = \rho(X_C)$ which ends up in the unauthorized copy. The strategy ρ can be deterministic or stochastic. The content owner uses an accusation algorithm σ. The output of the algorithm is a list of accused users.

The matrix X is constructed in two phases. In the first phase a list of random numbers $p = \{p_i\}_{i=1}^m$ is generated, where $p_i \in [t, 1-t]$, with t a small parameter satisfying $c_0 t \ll 1$. The p_i are independent and identically distributed according to a probability distribution function f. The function f is symmetric around $p_i = \frac{1}{2}$ and heavily biased towards values of p_i close to 0 and 1. In the second phase, the columns of X are filled by independently drawing random numbers $X_{ji} \in \{0,1\}$ with $\mathbb{P}[X_{ji} = 1] = p_i$.

Having spotted a copy with embedded mark y, the content owner computes an 'accusation sum' S_j for each user j according to

$$S_j = \sum_{i=1}^m y_i U(X_{ji}, p_i) \quad ; \quad U(X_{ji}, p_i) = \begin{cases} g_1(p_i) & \text{if } X_{ji} = 1 \\ g_0(p_i) & \text{if } X_{ji} = 0 \end{cases} \tag{1}$$

where g_0 and g_1 are the 'accusation functions', and y_i denotes the i'th bit of y. The decision whether to accuse a user is taken as follows: if $S_j > Z$, then accuse user j. Hence

$$\sigma(X, y, p) = \{j \mid S_j > Z\}. \tag{2}$$

We note some properties of this construction. If the received mark y_i is zero, the accusation of user j due to column i is neutral. If $y_i = 1$, the accusation sum is updated by $g_{X_{ji}}(p_i)$, which is a measure of how much suspicion arises from observing y_i for a given X_{ji} and p_i. The function $g_1(p)$ is positive and monotonically decreasing. The fact that user j has received a '1' in that position, i.e. $X_{ji} = 1$, adds to the suspicion. Moreover, the amount of suspicion decreases with increasing p_i, as the symbol becomes more probable. The function $g_0(p)$ is negative and monotonically decreasing. Therefore, the fact that $X_{ji} = 0$ detracts from the suspicion, and this becomes more pronounced for large p_i. We impose two properties on the accusation functions g_0 and g_1, which become handy during our proofs. First, we want to have the expectation value of the accusation in each column to be zero. Therefore, the functions should satisfy $pg_1(p) + (1-p)g_0(p) = 0$. Furthermore, we want the weights of the accusation for $X = 0$ and $X = 1$ to be symmetric, since the function $f(p)$ also has $0 \leftrightarrow 1$ symmetry. This is achieved by setting $g_0(p) = -g_1(1-p)$. These two properties together imply that (a) g_0 can be computed from g_1 according to $g_0(p) = -g_1(p) \cdot p/(1-p)$ and (b) g_1 on the interval $(\frac{1}{2}, 1-t)$ can be derived from g_1 on the interval $(t, \frac{1}{2})$ according to $g_1(p) = g_1(1-p) \cdot (1-p)/p$. Hence it is necessary only to specify $g_1(p)$ for $p \in (t, \frac{1}{2})$. [1]

The FP parameter ε_1, chosen by the content owner, denotes the desired bound on the probability of having $j \in \sigma(X, y, p)$ when a fixed user j is innocent. The FN parameter ε_2, also chosen by the content owner, denotes the desired bound on the probability that $\sigma(X, y, p)$ does not contain any guilty user.

The Tardos fingerprinting scheme uses a code length m and a threshold Z with the following scaling behavior[2] as a function of n, ε_1 and c_0:

$$m = Ac_0^2 \lceil \ln 1/\varepsilon_1 \rceil \quad ; \quad Z = Bc_0 \lceil \ln 1/\varepsilon_1 \rceil. \tag{3}$$

Here we have introduced the parameters A and B, replacing Tardos' constants 100 and 20, respectively. We want the scheme to satisfy the following properties.

---

[1] The constraint $g_1'(p) < 0$ on the whole interval $(t, 1-t)$ gives the additional condition $g_1'(p) > -\frac{g_1(p)}{p(1-p)}$. This can be satisfied by writing $g_1(p) = \exp \int^p \mathrm{d}p' \, Q(p')$ with $-\frac{1}{p(1-p)} < Q(p) < 0$.

[2] Note that (3) has no explicit ε_2-dependence. The dependence enters implicitly through A and B. However, as will be shown in Section 4, the ε_2-dependence vanishes in the limit of large c_0.

Definition 1: 'Soundness'

Let $\varepsilon_1 \in (0,1)$ be a fixed constant and let j be an arbitrary innocent user. We say that the above described fingerprinting scheme is ε_1-sound if, for all coalitions $C \subseteq [n]\backslash\{j\}$, and for all C-strategies ρ,

$$\mathbb{P}[\text{False positive}] = \mathbb{P}[j \in \sigma] < \varepsilon_1. \tag{4}$$

Definition 2: 'Completeness'

Let $\varepsilon_2 \in (0,1)$ and $c_0 \in \mathbb{N}^+$ be fixed constants. We say that the fingerprinting scheme is (c_0, ε_2)-complete if, for all coalitions C of size $c \leq c_0$, and all C-strategies ρ,

$$\mathbb{P}[\text{False negative}] = \mathbb{P}[C \cap \sigma = \emptyset] < \varepsilon_2. \tag{5}$$

Tardos proved (for $c_0 \geq 7$) that his scheme is Sound and Complete for the following very specific choice of parameters:[3]

$$f^{\mathrm{T}}(p) = \frac{1}{\pi - 4t'} \frac{1}{\sqrt{p(1-p)}}; \qquad g_1^{\mathrm{T}}(p) = \sqrt{\frac{1-p}{p}}; \qquad g_0^{\mathrm{T}}(p) = -\sqrt{\frac{p}{1-p}}$$

$$A^{\mathrm{T}} = 100; \qquad B^{\mathrm{T}} = 20; \qquad t^{\mathrm{T}} = \frac{1}{300c_0}$$

$$\varepsilon_2^{\mathrm{T}} = \varepsilon_1^{c_0/4} \tag{6}$$

where $t' = \arcsin\sqrt{t}$.

3 Proving a shorter code length through parametrization

3.1 Main result

The main aim of this paper is to show that it is possible to satisfy Soundness and Completeness in Tardos' scheme also with different choices, especially with a smaller parameter A and hence shorter code length m. While the choice $A^{\mathrm{T}} = 100$ does the job for $c_0 \geq 7$, we will show that it can be reduced to a number slightly larger than $4\pi^2$ in the limit of large c_0, when ε_2 is *not* coupled to ε_1. Our results can be summarized in the form of the following theorem.

Theorem 1: Let $\varepsilon_1, \varepsilon_2 \in (0,1)$ be fixed parameters. Let the cutoff parameter t be parametrized as $t = \tau/c_0$, with $0 < \tau \ll 1$. Let c_0 satisfy

$$c_0 \geq \frac{1}{\tau(3.4\pi)^2}. \tag{7}$$

Let $\omega \ll 1$ be a positive constant. Let the quantities D, δ and ξ be defined as

$$
\begin{aligned}
D &:= e(\omega/1.7)^2 \\
\delta &:= 2\tau + \pi\omega + e^{1.7}\frac{\pi c_0}{\omega(1-D)}D^{\tau + \frac{1.7\sqrt{\tau}}{\omega}\sqrt{c_0 - \tau}} \\
\xi &:= \sqrt{1 + \frac{1-\delta}{\pi\omega c_0}\cdot\frac{\ln\varepsilon_2}{\ln\varepsilon_1}} - 1.
\end{aligned}
\tag{8}
$$

Let the length m and the threshold Z in the generalized Tardos scheme be parametrized according to (3). Then it is possible to find functions $f(p)$, $g_1(p)$ such that the parameter setting

$$A = 4\pi^2\frac{(1+\xi)^2}{(1-\delta)^2} \quad ; \quad B = 4\pi\frac{1+\xi}{1-\delta}. \tag{9}$$

achieves ε_1-soundness and (c_0, ε_2)-completeness.

---

[3] In [Tar2003] the distribution f was given in terms of a uniform random variable r, defined according to $p = \sin^2 r$.

Corollary 1: In the limit of large c_0, ε_1-soundness and (c_0, ε_2)-completeness can be achieved by a code of length

$$m = 4\pi^2 \frac{1}{(1 - 2\tau - \pi\omega)^2}[1 + \mathcal{O}(\frac{1}{c_0})] \cdot c_0^2 \ln \frac{1}{\varepsilon_1}. \tag{10}$$

The proof of Theorem 1 is given in the coming sections; we follow Tardos' proof method wherever possible. We derive the conditions for Soundness and Completeness in Sections 3.2 and 3.3, respectively. We combine these conditions in Section 3.4 and obtain the lowest possible value of A that allows for Soundness and Completeness, depending on the specific choice of the functions f and g_1. In Section 3.5, we prove that $g_1 = g_1^{\mathrm{T}}$ is the 'optimal' choice, independent of f, in the sense that this choice minimizes this lowest value of A (given the proof method). In Section 3.6, we argue that, given $g_1 = g_1^{\mathrm{T}}$, the choice $f = f^{\mathrm{T}}$ is 'optimal' (in the same sense) within a limited class of functions. Finally, in Section 3.7, we set $f = f^{\mathrm{T}}$ and $g_1 = g_1^{\mathrm{T}}$ and complete the last step of the proof.

3.2 Condition for Soundness

This section follows the lines of Tardos' proof of Theorem 1 in [Tar2003]. We derive an inequality for the scheme parameters from the requirement that the Soundness property holds.

First, an auxiliary variable $\alpha_1 > 0$ is introduced for the purpose of applying the Markov inequality. (Tardos uses $\alpha_1^{\mathrm{T}} = 1/(10c_0)$.)

$$\mathbb{P}[j \in \sigma] = \mathbb{P}[S_j > Z] = \mathbb{P}[e^{\alpha_1 S_j} > e^{\alpha_1 Z}] < \frac{\mathbb{E}_{yXp}[e^{\alpha_1 S_j}]}{e^{\alpha_1 Z}}. \tag{11}$$

Here the notation $\mathbb{E}_{yXp}$ denotes the expectation value computed by averaging over all stochastic degrees of freedom: the (possibly stochastic) y_i, all the entries in X and the parameters p_i. The expectation value in (11) is bounded by using the inequality[4] $e^u \leq 1 + u + u^2$, which holds for $u < 1.7$. Using the notation $u_i = g_{X_{ji}}(p_i)$, we can write

$$\mathbb{E}_{X_j}[e^{\alpha_1 S_j}] = \prod_{i:y_i=1} \mathbb{E}_{X_{ji}}[e^{\alpha_1 u_i}] \leq \prod_{i:y_i=1} \left(1 + \alpha_1 \mathbb{E}_{X_{ji}}[u_i] + \alpha_1^2 \mathbb{E}_{X_{ji}}[u_i^2]\right) \tag{12}$$

Here $\mathbb{E}_{X_j}$ stands for the expectation value over the j'th row of X, keeping the rest of X fixed, and keeping $\{y_i\}$ and $\{p_i\}$ fixed. Note that y is independent of X_j since the user j is innocent. The inequality (12) holds as long as α_1 is so small that $\alpha_1 u_i < 1.7$ for all i for which $y_i = 1$. This is automatically true for those columns where $X_{ji} = 0$, since g_0 is negative. In the other columns, we need $\alpha_1 g_1(p_i) < 1.7$. Since $p_i \geq t$ and g_1 is monotonously decreasing, we can satisfy the inequality for all X by setting $\alpha_1 < \alpha_1^{\max} := 1.7/g_1(t)$.

Next we further bound (12). Due to the property $pg_1(p) + (1-p)g_0(p) = 0$ we have $\mathbb{E}_{X_{ji}}[u_i] = 0$. Thus we can write

$$\mathbb{E}_{X_j}[e^{\alpha_1 S_j}] \leq \prod_{i:y_i=1} \left(1 + \alpha_1^2 \mathbb{E}_{X_{ji}}[u_i^2]\right) \leq \prod_{i=1}^{m} \left(1 + \alpha_1^2 \mathbb{E}_{X_{ji}}[u_i^2]\right). \tag{13}$$

Note that all dependence on the coalition strategy has disappeared in the last expression.

Next we take the expectation value of (13), for fixed $\{p_i\}$, over the remaining degrees of freedom in X (all rows except j) and over $\{y_i\}_{i=1}^m$. This has no effect on the last expression in (13). Finally we take the expectation value $\mathbb{E}_p$ w.r.t. the p_i degrees of freedom. We remind the reader that this amounts to multiplying with the distribution function $\prod_{i=1}^m f(p_i)$ and integrating over all p_i. For ease of notation, we introduce the functional ν, defined as

$$\nu := \mathbb{E}_p\left[\mathbb{E}_{X_{ji}}[u_i^2]\right] = \int_t^{1-t} \mathrm{d}p\, f(p)\left\{p[g_1(p)]^2 + (1-p)[g_0(p)]^2\right\} = 2\int_t^{1/2} \mathrm{d}p\, f(p)\frac{p}{1-p}[g_1(p)]^2. \tag{14}$$

---

[4] One may ask why this crude inequality is used when our purpose is to squeeze the scheme's parameters as tightly as possible. The answer is that we do not want to deviate from [Tar2003] too much in this paper, as it would lead to even more bookkeeping than is already the case. It would be interesting to determine the consequences of using an inequality of the more general form $e^u < 1 + ru + su^2$, with $r \geq 1$.

(With Tardos' choice for g_1, this evaluates to $\nu = \int_t^{1-t} \mathrm{d}p\, f(p) = 1$). We can then write

$$\mathbb{E}_{yXp}[e^{\alpha_1 S_j}] \leq \prod_{i=1}^m (1 + \nu \alpha_1^2) \leq e^{m\nu \alpha_1^2}. \tag{15}$$

In the last step we have used $1 + u \leq e^u$, which holds for all u. Substitution of (15) into (11) finally gives

$$\mathbb{P}[S_j > Z] < \exp(m\nu \alpha_1^2 - \alpha_1 Z). \tag{16}$$

As (16) holds for all α_1 in the allowed range, we can write

$$\mathbb{P}[S_j > Z] < \min_{\alpha_1 \in (0, \alpha_1^{\max})} \exp(m\nu \alpha_1^2 - \alpha_1 Z). \tag{17}$$

The minimum of the parabola in the exponent lies at $\alpha_1^* := Z/(2m\nu)$. Hence, the minimum in (17) is obtained by setting $\alpha_1 = \alpha_1^*$. Note that this is allowed only if $\alpha_1^{\max} \geq \alpha_1^*$; this condition can be rewritten as

$$c_0 \geq \frac{Bg_1(t)}{3.4\nu A}. \tag{18}$$

If c_0 is large enough for the condition to be satisfied, then it holds that

$$\mathbb{P}[S_j > Z] < \exp(m\nu [\alpha_1^*]^2 - \alpha_1^* Z) < \varepsilon_1^{B^2/(4\nu A)}. \tag{19}$$

In the last inequality we have used $\alpha_1^* = Z/(2m\nu)$ and the parametrization (3). From (19) and Definition 1 we conclude that, for c_0 large enough so that (18) holds, ε_1-soundness can only be obtained if

$$A \leq \frac{B^2}{4\nu}. \tag{20}$$

3.3 Condition for Completeness

This section closely follows the proof of Theorem 2 in [Tar2003]. We derive an inequality for the scheme parameters from the Completeness requirement. First, the coalition's accusation sum S is defined,

$$S = \sum_{j \in C} S_j = \sum_{i=1}^m y_i \{x_i g_1(p_i) + [c - x_i] g_0(p_i)\}. \tag{21}$$

Here $x_i = \sum_{j \in C} X_{ji}$ denotes the number of colluders that have a '1' at the i-th position of their codeword. (The size of the coalition is $c = |C|$. We consider the case $c \leq c_0$.) Since $S > cZ$ would imply that at least one colluder gets accused, the false negative error probability can be bounded by

$$\mathbb{P}[C \cap \sigma = \emptyset] \leq \mathbb{P}[S \leq cZ]. \tag{22}$$

An auxiliary constant $\alpha_2 > 0$ is introduced for the purpose of using the Markov inequality (Tardos chooses $\alpha_2^{\mathrm{T}} = 1/(20c_0)$),

$$\mathbb{P}[S \leq cZ] = \mathbb{P}[e^{-\alpha_2 S} \geq e^{-\alpha_2 cZ}] \leq \frac{\mathbb{E}_{yXp}[e^{-\alpha_2 S}]}{e^{-\alpha_2 cZ}}. \tag{23}$$

Upper bounding the expectation value $\mathbb{E}_{yXp}[e^{-\alpha_2 S}]$ is an arduous job. The derivation is given in Appendix A. It turns out that $\mathbb{E}_{yXp}[e^{-\alpha_2 S}] < \exp(-\alpha_2 m/L)$, with $L > 0$ a numerical constant, provided that c, L, t and α_2 satisfy a complicated condition (88). (Tardos has $L^{\mathrm{T}} = 4$). Thus we can bound the false negative probability as

$$\mathbb{P}[C \cap \sigma = \emptyset] < \exp(\alpha_2 cZ - \alpha_2 m/L) \leq \exp(\alpha_2 c_0 Z - \alpha_2 m/L). \tag{24}$$

Substituting the parametrization (3) into (24), and restricting ourselves to the regime $A > LB$, we get

$$\mathbb{P}[C \cap \sigma = \emptyset] < \varepsilon_1^{\alpha_2 c_0^2 [A/L - B]}. \tag{25}$$

Next we demand Completeness, $\mathbb{P}[C \cap \sigma = \emptyset] < \varepsilon_2$. This yields the following inequality,

$$A \geq LB + \frac{L}{\alpha_2 c_0^2} \frac{\ln \varepsilon_2}{\ln \varepsilon_1}. \tag{26}$$

Note that (26) is valid only if the complicated condition (88) is satisfied.

3.4   Conditions on the code length

By combining the results of Sections 3.2 and 3.3, we derive a condition on the length parameter A that is sufficient for proving Soundness and Completeness using the current proof technique.

Lemma 1: Let $\varepsilon_1, \varepsilon_2 \in (0, 1)$ be fixed parameters. Let the code length m and the threshold Z in the generalized Tardos scheme be parametrized in terms of the A and B parameters according to (3). Let the threshold t be set as $t = \tau/c_0$, with $\tau \ll 1$ a positive constant. Let c_0 be a fixed integer, satisfying

$$c_0 \geq \frac{1}{\tau(3.4\nu\pi)^2}. \tag{27}$$

Let α_2 be the auxiliary parameter introduced in Section 3.3. Let L be the parameter introduced in Section 3.3, chosen such that the condition (88) can be satisfied by some value α_2. Let ν be the functional of f and g_1 as defined by (14). Let ψ be defined as

$$\psi = \sqrt{1 + \frac{1}{\nu L \alpha_2 c_0^2} \frac{\ln \varepsilon_2}{\ln \varepsilon_1}} - 1. \tag{28}$$

Then the generalized Tardos scheme with

$$A = 4\nu L^2 (1 + \psi)^2 \quad ; \quad B = 4\nu L(1 + \psi) \tag{29}$$

is ε_1-sound and (c_0, ε_2)-complete.

Corollary 2: Let α_2 be parametrized as $\alpha_2 = \omega/c_0$, with $\omega < 1$ a positive constant. Let the function $g_1(p)$ be such that it has the asymptotic behavior $g_1(p) \propto p^{-\gamma}$ at $p = t$, with $\gamma < 1$. Then for large c_0, it is possible to achieve Soundness and Completeness with $A = 4\nu L^2[1 + \mathcal{O}(1/c_0)]$.

Proof of Lemma 1: The value of c_0 in (27) is specifically set such that (18) is satisfied, and hence we are allowed to use the inequality (20). By combining the results (20) and (26) we obtain a 'window' for A in which Soundness and Completeness are both satisfied,

$$A \in \left[LB + \frac{L}{\alpha_2 c_0^2} \frac{\ln \varepsilon_2}{\ln \varepsilon_1}, \quad \frac{B^2}{4\nu} \right]. \tag{30}$$

We choose B such that (i) this window exists, and (ii) the left boundary is as small as possible. It turns out that the optimal choice for B is the smallest value for which the window still exists. Setting the left and right boundary in (30) equal to each other and solving for B gives

$$B_{\text{opt}} = 4\nu L(1 + \psi), \tag{31}$$

with ψ as defined in (28). The corresponding value for A is $A_{\text{opt}} = B_{\text{opt}}^2/(4\nu) = 4\nu L^2(1 + \psi)^2$. $\square$

Proof of Corollary 2: First we show that the quantity $L > 0$ is well defined in the limit of large c_0. To this end we inspect (88), setting $c = c_0$. We start with the term involving $\Delta = e[\omega/1.7]^{1/(1-\beta)} < 1$, (with $\omega < 1$ and $\frac{1}{1-\beta} > 1$). The exponent $c_0 - x_{\max}$, with $x_{\max}$ defined in (82), satisfies

$$c_0 - x_{\max} \geq \tau + \frac{1.7(1 - t)c_0}{\omega g_1(\tau/c_0)}. \tag{32}$$

From the choice $g_1(t) \propto t^{-\gamma}$, with $\gamma < 1$, it follows that $c_0 - x_{\max} = \mathcal{O}(c_0^{1-\gamma})$, i.e. an increasing function of c_0. Hence the expression $\Delta^{c_0-x_{\max}}/(1-\Delta)$ decreases as $\exp[-\lambda c_0^{1-\gamma}]$ for some positive constant λ. For large c_0, this is much faster than the term $\alpha_2/L = \omega/c_0 L$.

Next we look at the term $\nu c_0 \alpha_2^2$. This can be written as $\alpha_2 \nu \omega$. Hence, the factor multiplying α_2 is a small constant independent of c_0.

In the term $\alpha_2 c_0 \mathbb{E}_p[p^{c_0} g_1(p)]$, the factor multiplying α_2 can be written in the form

$$c_0 \mathbb{E}_p[p^{c_0} g_1(p)] = [pf(p)g_1(p)p^{c_0}]_{p=\tau/c_0}^{1-\tau/c_0} - \int_{\tau/c_0}^{1-\tau/c_0} \mathrm{d}p\, p^{c_0} \frac{\mathrm{d}}{\mathrm{d}p}[pf(p)g_1(p)]. \tag{33}$$

In this form, it is clear that the expression is finite for $c_0 \to \infty$, provided that the product $f(p)g_1(p)$ is non-pathological near $p = 1 - t$.

The remaining term in (88) is $\alpha_2 \cdot \sum_{x=1}^{c_0-1} \binom{c_0}{x} K_x^{\mathrm{bound}}$, where K_x^{bound} is some nonnegative number that upper bounds the expression $K_{1,x}$ as defined in (75). From (75) it follows that

$$K_{1,x} \leq tf(t)g_1(t)t^x(1-t)^{c_0-x} + \int_J \mathrm{d}p\, p^x(1-p)^{c_0-x}\frac{\mathrm{d}}{\mathrm{d}p}(pfg_1), \tag{34}$$

where J is defined as the interval (or set of intervals) on which $\frac{\mathrm{d}}{\mathrm{d}p}(pfg_1) > 0$. Using the expression (34) as our bound K_x^{bound} and computing the x-sum we obtain

$$\sum_{x=1}^{c_0-1} \binom{c_0}{x} K_x^{\mathrm{bound}} = tf(t)g_1(t)[1 - (1-t)^{c_0} - t^{c_0}] + \int_J \mathrm{d}p\, [1 - (1-p)^{c_0} - p^{c_0}]\frac{\mathrm{d}}{\mathrm{d}p}(pfg_1). \tag{35}$$

Expressed in this form, it is clear that this contribution is of order $1 - (1-t)^{c_0} = \mathcal{O}(\tau)$ for $c_0 \gg 1$. The smallness of all the expressions in (88) implies that L is well-defined.

The last step in the proof of Corollary 2 is the asymptotic behavior of the factor $(1 + \psi)^2$ in (29). The fraction $1/(\alpha_2 c_0^2)$ in the definition (28) of ψ is equal to $1/(\omega c_0)$, i.e. of order $\mathcal{O}(1/c_0)$. Consequently, $(1 + \psi)^2 = 1 + \mathcal{O}(1/c_0)$. □

Note that the expression $4\nu L^2$ in Lemma 1 originates from the specific proof technique.

The next step in the proof of Theorem 1 is to choose f and g_1 such that the product νL^2 is minimized. We refer to this choice as 'optimal', but it should be clear that it represents optimality only with respect to the proof technique that we employ in Section 3.

3.5   Finding the optimal g_1 function

Lemma 2: For all distributions f, Tardos' choice $g_1 = g_1^{\mathrm{T}}$ is optimal in the sense that it minimizes the factor νL^2 appearing in the length parameter A in Lemma 1. This choice yields $\nu = 1$.

Proof: We have to solve an optimization problem and determine where the functional derivative of $4\nu L^2$ is zero. This is easiest to accomplish by choosing as the independent degrees of freedom $f(p)$ and $s(p) := pf(p)g_1(p)$ on the interval $p \in (t, 1/2)$, instead of f and g_1.

From (88) it can be seen that L depends on f and g_1 only through $s(p)$: The expression $\mathbb{E}_p[p^c g_1]$ involves only the product fg_1; and K_x^{bound} is a bound on (75), which also depends on f and g_1 solely through the product fg_1.

The parameter ν, on the other hand, depends on both the $s(p)$ and $f(p)$ degrees of freedom; Eq.(14) can be rewritten as

$$\nu[s, f] = 2 \int_t^{1/2} \mathrm{d}p\, \frac{s^2(p)}{p(1-p)f(p)}. \tag{36}$$

The functional that we have to minimize is

$$\mathcal{F}[s, f] = 4\nu[s, f]L^2[s] + \lambda \left[\int_t^{1/2} \mathrm{d}p\, f(p) - 1/2\right], \tag{37}$$

where λ is a Lagrange multiplier for the normalization constraint on f. Setting the functional derivative with respect to $f(p)$ equal to zero gives

$$0 = \frac{\delta \mathcal{F}}{\delta f(p)} = -8L^2 \frac{s^2(p)}{p(1-p)f^2(p)} + \lambda \quad \implies \quad g_1(p) \propto \sqrt{\frac{1-p}{p}}. \tag{38}$$

Hence, since the normalization of g_1 is arbitrary[5], it turns out that Tardos' choice $g_1(p) = g_1^{\mathrm{T}}(p) = \sqrt{(1-p)/p}$ is optimal.

From the optimal g_1 the value of ν follows directly, without dependence on f: Substituting g_1^{T} into (14) and using the fact that f is a normalized probability distribution, we get $\nu = 1$. □

3.6   The choice $f = f^{\mathrm{T}}$ seems to be optimal

Having found the optimal g_1, we can next search for the optimal distribution function f. The terms proportional to α_2 in the left hand side of (88) are the most important in determining the allowed values of L: we want to tune f such that the number multiplying α_2 is as negative as possible. We have looked at a class of smooth functions of the form $f(p) = \frac{1}{2}p^{a-1}(1-p)^{b-1}/[B_{1/2}(a,b)-B_t(a,b)]$ on the interval $p \in (t, 1/2)$, where B is the incomplete Beta function. Numerical inspection shows that Tardos' choice $a = b = \frac{1}{2}$ is the best choice for a, b. Of course, as we have not investigated the full function space of f, this does not prove that f^{T} is the best possible choice.

In the rest of this paper we will work with $f = f^{\mathrm{T}}$.

3.7   Last step in the proof of Theorem 1

We are now finally in a position to prove Theorem 1. We show that Theorem 1 follows from Lemma 1 in the special case $f = f^{\mathrm{T}}$ and $g_1 = g_1^{\mathrm{T}}$. In this special case, we have $\nu = 1$ as shown by Lemma 2. Hence the condition on c_0 (27) reduces to (7). Furthermore, the choice $g_1 = g_1^{\mathrm{T}}$ gives $\beta = \frac{1}{2}$ (see App. A), whereby Δ reduces to D, as defined in (8). Next we use the property $\frac{\mathrm{d}}{\mathrm{d}p}[pf^{\mathrm{T}}g_1^{\mathrm{T}}] = 0$ to explicitly evaluate the $\mathbb{E}_p$ expectation and the K_x^{bound} term in (88). We make use of expression (35) for K_x^{bound} and get

$$c_0\mathbb{E}_p[p^{c_0}g_1(p)] = \frac{(1-t)^{c_0} - t^{c_0}}{\pi - 4t'} \quad ; \quad \sum_{x=1}^{c_0-1} \binom{c_0}{x} K_x^{\mathrm{bound}} = \frac{1 - (1-t)^{c_0} - t^{c_0}}{\pi - 4t'}. \tag{39}$$

This allows us to rewrite (88) as

$$1 - \alpha_2 \left\{ \frac{2(1-t)^{c_0} - 1}{\pi - 4t'} - \omega - e^{1.7} \frac{c_0}{\omega(1-D)} D^{\tau + 1.7\omega^{-1}\sqrt{\tau}\sqrt{c_0 - \tau}} \right\} < 1 - \frac{\alpha_2}{L}. \tag{40}$$

The quantity δ in Theorem 1 is specifically chosen such that $L = \pi/(1-\delta)$ satisfies (40). The ξ in Theorem 1 is given by ψ (28) after the substitution $\nu = 1$, $\alpha_2 c_0 = \omega$, $L = \pi/(1-\delta)$. Thus, (29) reduces to (9). □

4   Numerical evaluation

In the preceding section, we proved a result for large coalition sizes. In this section, we numerically investigate how quickly convergence to this behavior occurs.

4.1   Method

Given that f^{T} and g_1^{T} are the optimal functions, we determine the optimal values for the remaining parameters. For fixed $(c_0, \varepsilon_1, \varepsilon_2)$, our task is to find $(t, \alpha_1, \alpha_2, L, B)$ such that we obtain the smallest possible A. We mean 'smallest' in the sense that Soundness and Completeness can be

---

[5] As can be seen from the accusation rule (2), rescaling g_1 and Z by the same factor leaves the scheme invariant.

proven using the technique in Section 3, but without the assumption that c_0 is large. The following constraints must be satisfied:

- From (16), (3) and the required property $\mathbb{P}[\text{false positive}] < \varepsilon_1$ we get the constraint

$$- Ac_0^2\alpha_1^2 + Bc_0\alpha_1 < 1. \tag{41}$$

- Similarly, from (24), (3) and the required property $\mathbb{P}[\text{false negative}] < \varepsilon_2$ we get the constraint

$$(-B + A/L)c_0^2\alpha_2 \ln \varepsilon_1 < \ln \varepsilon_2. \tag{42}$$

- In Section 3.2 the parameter α_1 was introduced such that $\alpha_1 < 1.7/g_1(t)$. This gives

$$\alpha_1 < 1.7\sqrt{t/(1-t)}. \tag{43}$$

- In Appendix A, it is assumed that α_2 is so small that the function $(1-p)^{c-x}\exp[\alpha_2(cp - x)/\sqrt{p(1-p)}]$ is a decreasing function of p in the vicinity of $p = 1-t$ for all $x \in \{1, \ldots, c-1\}$. Let's denote this function as $z(p)$. Its derivative near $p = 1-t$ is approximately given by $\mathrm{d}z/\mathrm{d}p \approx -z \cdot (c-x)(1-p)^{-1}[1 - \frac{\alpha_2}{2\sqrt{1-p}}]$. Hence, in order to ensure a negative sign of the derivative, α_2 has to satisfy

$$\alpha_2 \lesssim 2\sqrt{t}. \tag{44}$$

- In condition (88) with $f = f^{\mathrm{T}}$, $g_1 = g_1^{\mathrm{T}}$ the K_x^{bound} term is easily evaluated. The resulting condition is

$$1 - \alpha_2\frac{2(1-t)^{c_0} - 1}{\pi - 4t'} + c_0\alpha_2^2 + e^{1.7}\frac{D^{c_0-x_{\max}}}{1-D} < 1 - \frac{\alpha_2}{L} \tag{45}$$

with $D = e(c_0\alpha_2/1.7)^2$ and $c_0 - x_{\max} = \lceil c_0 t + 1.7\sqrt{t(1-t)}/\alpha_2 \rceil \geq 1$.

The complicated α_2-dependence of (45), containing a term of the form α_2^{1/α_2}, prevents us from finding an optimum analytically. Instead, we have searched for optimum parameter values numerically, using a randomized method following these steps:

1. Choose a random t uniformly from the interval $(0, \frac{1}{2c_0})$.

2. Choose α_1 uniformly from $(0, \min\left\{1.7\sqrt{t/(1-t)}, \frac{1}{c_0}\frac{2(1-t)^{c_0}-1}{\pi-4t'}\right\})$.

3. Choose L uniformly from $(\frac{\pi-4t'}{2(1-t)^{c_0}-1}, \frac{1}{c_0\alpha_1})$.

4. Find the largest possible value of α_2 in the interval $\left(0, \min\left\{2\sqrt{t}, \frac{1}{c_0}\left[\frac{2(1-t)^{c_0}-1}{\pi-4t'} - \frac{1}{L}\right]\right\}\right)$ that satisfies condition (45).

5. Compute $A = \frac{L(c_0\alpha_1)^{-1}}{(c_0\alpha_1)^{-1}-L}\left[(c_0\alpha_1)^{-1} + \frac{1}{c_0^2\alpha_2}\frac{\ln \varepsilon_2}{\ln \varepsilon_1}\right]$.

Note that the optimal choice of B (in terms of achieving small A) follows from (30). When the optimal value for B is used, the interval (30) consists of a single point seen in the last step. We repeat steps 1–5 multiple times and select the set of parameter values that yield the lowest value of A, i.e. the shortest code length.

4.2   Numerical results

As the parameter A depends on ε_1 and ε_2 only through the ratio of their logarithms, we define the parameter $R = \frac{\ln \varepsilon_2}{\ln \varepsilon_1}$ to display our results. This allows us to represent dependence on three parameters using only two variables, namely $A_{\text{best}} = A_{\text{best}}(c_0, R)$. As we are primarily concerned with content distribution applications, we have chosen $10^{-15} \leq \varepsilon_1 \leq 10^{-9}$ and $0.1 \leq \varepsilon_2 \leq 0.5$ as plausible values. This gives $0.02 \lessapprox R \lessapprox 0.1$. We further consider $10 \leq c_0 \leq 80$.

We plot the best code length parameter A as a function of c_0 (for constant R) and as a function of R (for constant c_0) in Fig. 1 and Fig. 2, respectively. Note that these figures are derived from the same dataset. In Appendix B, we give the corresponding values of the t and B parameters, which are necessary to implement the fingerprinting scheme. The numerical results indicate that the result $A \approx 4\pi^2 \approx 39.5$ (see Corollary 1) seems to be approached for large c_0 and small R. Moreover, The R-dependence of A is slightly sublinear. Note that there is not much variation in the value of A, only of the order of 15%.

In Fig. 3, we also plot the results when the false positive and false negative rates are coupled as in [Tar2003], i.e. $\varepsilon_2 = \varepsilon_1^{c_0/4}$ or $R = c_0/4$. In this case, it is possible to reduce A to approximately 90, which is not much of an improvement with respect to Tardos' $A^{\text{T}} = 100$. This result further emphasizes the importance of decoupling FP and FN rates. When we allow for high FN rates $\varepsilon_1 \ll \varepsilon_2$, the code can be safely made more than a factor two shorter than suggested in Tardos' original construction.

Finally, the figures in this section, together with Appendix B, give a system parameter recipe for content owners who wish to implement a provably ε_1-sound and (c_0, ε_2)-complete fingerprinting scheme with a code length $A(R, c_0) c_0^2 \ln \varepsilon_1^{-1}$.
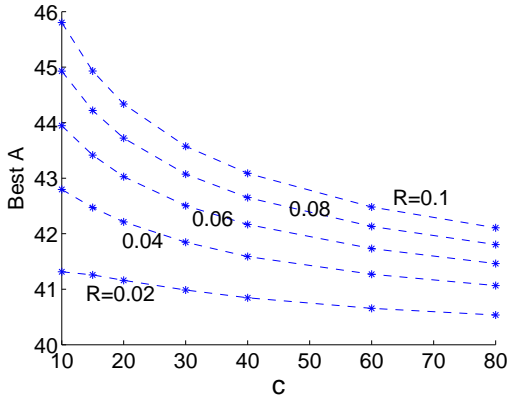

Fig. 1: *Smallest obtained A as a function of c_0 for various values of $R = \frac{\ln \varepsilon_2}{\ln \varepsilon_1}$.*
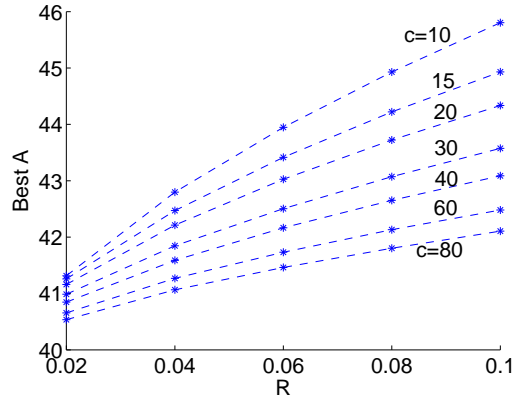
Fig. 2: *Smallest obtained A as a function of $R = \frac{\ln \varepsilon_2}{\ln \varepsilon_1}$ for various values of c_0.*

5   Statistical approach

5.1   Motivation and main result of the statistical approach

In this section, we put aside the provable properties of the Tardos fingerprinting scheme. Instead we study the statistical behavior of the accusation sums S_j (1) for the innocent user and S (21) for the coalition. The advantage of the statistical approach is that we get more insight into the 'true' behavior of the scheme (actual FP and FN probabilities as a function of m, c_0, n, Z) than provided by the provable result (Theorem 1) based on the Markov inequality.

The accusation sums S_j and S are defined as the sum of a large number of stochastic variables. We expect S_j and S to have an approximately Gaussian probability distribution. (This is motivated in Appendix C.) The Central Limit Theorem states that a variable which is created by
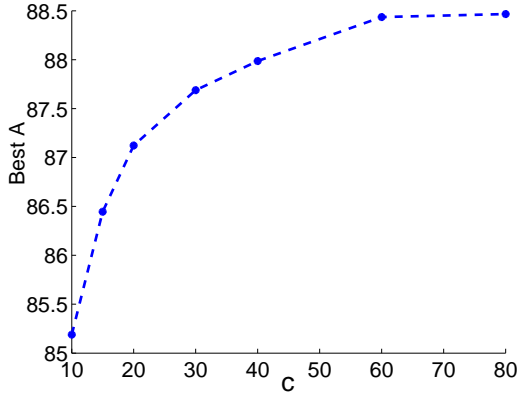
Fig. 3: *Smallest obtained A as a function of c_0*
in the coupled case $R = c_0/4$.

adding up many independent stochastic variables will have a Gaussian distribution in the vicinity of the distribution's peak. The exact size of this 'vicinity' depends on the distribution of the individual variables and on the total number of variables in the sum. In Appendix C, we argue that the Gaussian approximation is valid for some realistic values of ε_1 and m.

Our main result can be formulated as follows.

Theorem 2: Let $\varepsilon_1, \varepsilon_2 \in (0,1)$ be fixed parameters. Let $c_0 \in \mathbb{N}^+$ be a fixed parameter. Let the functions $f(p)$ and $g_1(p)$ be given by $f = f^{\mathrm{T}}$, $g_1 = g_1^{\mathrm{T}}$. Let the cutoff parameter t be parametrized as $t = \tau/c_0$. Let the accusation sums S_j (1) and S (21) obey Gaussian statistics. Then the fingerprinting scheme with code length m and threshold Z set according to

$$m \;\geq\; \frac{2\pi^2}{(1-2\tau)^2}\, c_0^2 \left[\mathrm{Erfc}^{\mathrm{inv}}(2\varepsilon_1) + \frac{1}{\sqrt{c_0}}\mathrm{Erfc}^{\mathrm{inv}}(2\varepsilon_2)\right]^2 \tag{46}$$

$$Z \;\in\; \left[\sqrt{2m}\,\mathrm{Erfc}^{\mathrm{inv}}(2\varepsilon_1), \quad \frac{1-2\tau}{\pi c_0}m - \frac{\sqrt{2m}}{\sqrt{c_0}}\mathrm{Erfc}^{\mathrm{inv}}(2\varepsilon_2)\right] \tag{47}$$

is ε_1-sound and (c_0, ε_2)-complete.

Here Erfc stands for the complementary error function $1 - \mathrm{Erf}$, with the definition $\mathrm{Erf}(x) = (2/\sqrt{\pi}) \int_0^x \mathrm{d}y\, e^{-y^2}$. The superscript 'inv' denotes the inverse function.

Corollary 3: Let $(\varepsilon_1, \varepsilon_2) \in (0,1)$ be independent fixed parameters. Then for $c_0 \gg 1$ the parameter choice

$$m = \frac{2\pi^2}{(1-2\tau)^2}c_0^2 \ln \frac{1}{\varepsilon_1\sqrt{2\pi}} \quad ; \quad Z = \frac{2\pi}{1-2\tau}c_0 \ln \frac{1}{\varepsilon_1\sqrt{2\pi}} \tag{48}$$

achieves ε_1-soundness and (c_0, ε_2)-completeness.

The proof of Theorem 2 is given in the coming sections and has the following outline. First, in Sections 5.2 and 5.3, we compute the lowest moments of the distributions of the accusation sums,

$$\mu_j := \mathbb{E}_{yXp}[S_j] \quad ; \quad \sigma_j^2 := \mathbb{E}_{yXp}[S_j^2] - \mu_j^2 \qquad \text{where } j \text{ is not a colluder}$$

$$\mu := \mathbb{E}_{yXp}[S] \quad ; \quad \sigma^2 := \mathbb{E}_{yXp}[S^2] - \mu^2. \tag{49}$$

Then, in Section 5.4, we compute the false positive and false negative error probabilities as a function of m, Z and c_0. We derive conditions on m and Z from the Soundness and Completeness

requirements. In Section 5.5, we identify an 'extremal' strategy which leads to a maximum value of σ_j and σ. In Section 5.6, we assume that the probability distributions are Gaussian (a motivation for this step is given in Appendix C) and we combine all the ingredients to complete the final step in the proof of Theorem 2. We prove Corollary 3 in Section 5.7.

5.2   Statistics of an innocent user's accusation

Even without knowing the colluders' strategy, we can derive a number of useful properties of the expectation values listed in (49). We start by looking at S_j, where user j is not a colluder. In Section 2.2, the functions g_1, g_0 were introduced such that $pg_1(p) + (1 - p)g_0(p) = 0$. This immediately yields

$$\mathbb{E}_{X_j}[S_j] = \sum_{i=1}^{m} y_i \, \mathbb{E}_{X_{ji}}[U(X_{ji}, p_i)] = 0, \tag{50}$$

where y_i is shifted out of the expectation value because j is not part of the coalition. From (50) it immediately follows that $\mu_j = 0$.

The standard deviation σ_j is computed as follows. Substitution of the definition (1) into (49) gives

$$\sigma_j^2 = \mathbb{E}_{yXp}[S_j^2] = \sum_{i=1}^{m} \sum_{k=1}^{m} \mathbb{E}_{yXp}[y_i y_k \, U(X_{ji}, p_i) U(X_{jk}, p_k)]. \tag{51}$$

All terms with $i \neq k$ vanish, since then the expectation value factorizes into two parts that are both zero due to (50). Hence we can write

$$\sigma_j^2 = \sum_{i=1}^{m} \mathbb{E}_{yXp}[y_i \, U^2(X_{ji}, p_i)] = \sum_{i=1}^{m} \mathbb{E}_{yX_Cp} \left[y_i \, \mathbb{E}_{X_{ji}}[U^2(X_{ji}, p_i)] \right]. \tag{52}$$

(Here the expectation $\mathbb{E}_{yX_Cp}$ involves only those entries in X that are visible to the colluders.) Again we have used the fact that y_i does not depend on X_{ji} when user j is innocent. Next we make use of the property $\mathbb{E}_{X_{ji}}[U^2(X_{ji}, p_i)] = 1$ which holds for $g_1 = g_1^{\mathrm{T}}$. This finally yields

$$\sigma_j^2 = \sum_{i=1}^{m} \mathbb{E}_{yX_Cp}[y_i] < m. \tag{53}$$

5.3   Statistics of the coalition accusation S

Next we look at the collective accusation sum S defined in (21). Now we have to keep in mind that y_i depends on X_{ji} when j is a colluder. Taking the expectation of (21) we get

$$\mu = \mathbb{E}_{yXp}[S] = \sum_{i=1}^{m} \mathbb{E}_p \left[\mathbb{E}_X \left[\mathbb{E}_y[y_i] \{ x_i g_1(p_i) + [c - x_i] g_0(p_i) \} \right] \right]. \tag{54}$$

The notation $\mathbb{E}_y$ stands for the expectation value over the y degrees of freedom for fixed p and X. In (54) the expectation value over X, for fixed p, reduces to a binomial distribution on the integers x_i [Tar2003]:

$$\mathbb{P}[\#'1' \text{ entries in column } i \text{ is } x_i] = \binom{c}{x_i} p_i^{x_i} (1 - p_i)^{c - x_i}. \tag{55}$$

We now evaluate (54) as follows. We express the $\mathbb{E}_X$ expectation in the form (55). We define a quantity $\psi_i(x_i) := \mathbb{E}_{X_Cp\backslash i}[\mathbb{E}_y[y_i]]$. Here the notation $\mathbb{E}_{X_Cp\backslash i}$ means the expectation value over all degrees of freedom in X_C and p except column i. The quantity ψ_i depends on x_i and the colluder strategy (and possibly explicitly on i, if the colluders choose to apply different strategies

in different positions); it does not depend on p_i, as the colluders do not have access to p_i. Finally we substitute Tardos' functions g_1^{T}, g_0^{T} and f^{T}. In this way we obtain, after some algebra,

$$\mu = \frac{1}{\pi - 4t'} \sum_{i=1}^{m} \sum_{x_i=0}^{c} \binom{c}{x_i} \psi_i(x_i) \left\{ (1-t)^{x_i} t^{c-x_i} - t^{x_i}(1-t)^{c-x_i} \right\}. \tag{56}$$

We now make use of the marking condition, giving $\psi_i(0) = 0$ and $\psi_i(c) = 1$. This allows us to rewrite (56) as

$$\mu = \frac{1}{\pi - 4t'} \sum_{i=1}^{m} \left[(1-t)^c - t^c + \sum_{x_i=1}^{c-1} \binom{c}{x_i} \psi_i(x_i) \left\{ (1-t)^{x_i} t^{c-x_i} - t^{x_i}(1-t)^{c-x_i} \right\} \right]. \tag{57}$$

Note that the coalition strategy has an almost negligible effect on μ. The terms $\binom{c}{x} t^x (1-t)^{c-x}$ add up to 1 when the full sum is taken, but only the $x = 0$ term is of order 1. All the other terms summed together are only of order $c_0 t = \tau \ll 1$. The same argument holds for the other summand, but there the $x = c$ term is dominant.

We use the same methods as above to evaluate σ. Without showing the details of the computation, we give the result,

$$\sigma^2 = \frac{1}{\pi - 4t'} \sum_{i=1}^{m} \sum_{x_i=1}^{c} \binom{c}{x_i} \psi_i(x_i) \int_t^{1-t} \mathrm{d}p \, (x_i - cp)^2 p^{x_i - 3/2} (1-p)^{c - x_i - 3/2} - \frac{\mu^2}{m}. \tag{58}$$

5.4 Relating m and Z to the error probabilities

In (53,57,58) we see from the i-summations that the m-dependence becomes very simple if the colluders apply the same strategy in each column; namely, the quantities σ_j^2, μ and σ^2 then all become proportional to m. This motivates us to define 'scaled' quantities as follows,

$$\sigma_j^2 = m\tilde{\sigma}_j^2 \qquad \mu = m\tilde{\mu} \qquad \sigma^2 = m\tilde{\sigma}^2. \tag{59}$$

Let us introduce the notation ρ_1 and ρ_2 for the probability distribution functions of S_j and S, respectively. These functions are unknown to us, but we normalize them so that they have zero mean and unit variance,

$$\mathbb{P}[S_j \in [s, s + \triangle s]] = \rho_1\left(\frac{s}{\sigma_j}\right)\frac{\triangle s}{\sigma_j} \quad ; \quad \mathbb{P}[S \in [s, s + \triangle s]] = \rho_2\left(\frac{s - \mu}{\sigma}\right)\frac{\triangle s}{\sigma}, \tag{60}$$

with $\int_{-\infty}^{\infty} \mathrm{d}x \, \rho_1(x) = 1$ and $\int_{-\infty}^{\infty} \mathrm{d}x \, \rho_2(x) = 1$. We introduce the cumulative 'tail' functions as

$$G_1(x) = \int_x^{\infty} \mathrm{d}x' \rho_1(x') \quad ; \quad G_2(x) = \int_{-\infty}^{x} \mathrm{d}x' \rho_2(x'). \tag{61}$$

With this notation, the error probabilities are expressed as

$$\text{FP error prob.} = G_1\left(\frac{Z}{\sigma_j}\right) \quad ; \quad \text{FN error prob.} = G_2\left(\frac{c_0 Z - \mu}{\sigma}\right). \tag{62}$$

This is sketched in Fig. 4. The left curve is the probability density of the quantity $S_j/\sqrt{m}$. It has mean $\tilde{\mu}_j = 0$ and variance $\tilde{\sigma}_j \approx 1$ (this will be shown in Section 5.5). The FP error rate (which should be less than ε_1) is given by the area to the right of the (rescaled) threshold $Z/\sqrt{m}$. The right curve is the probability density of the quantity $\frac{1}{c}S/\sqrt{m}$. It has average $\frac{1}{c}\tilde{\mu}\sqrt{m}$ and variance $\tilde{\sigma}/c$. The FN error rate is given by the area to the left of $Z/\sqrt{m}$. The horizontal axis is scaled such that the S_j-curve does not depend on c and m. Several important properties follow from this picture:
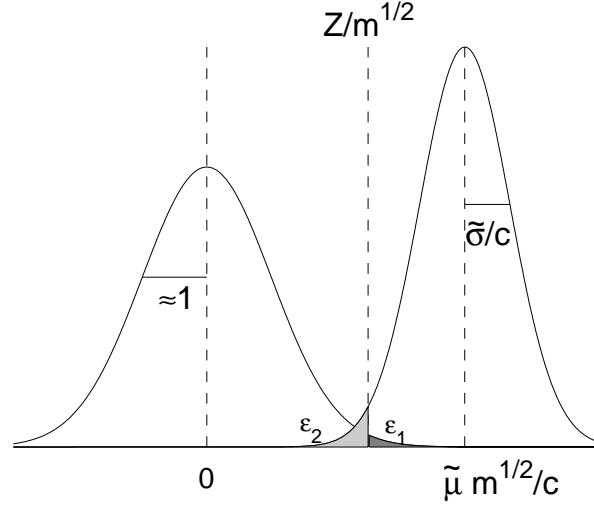
Fig. 4: Sketch of the probability density of $S_j/\sqrt{m}$ (left) and $\frac{1}{c}S/\sqrt{m}$ (right). The accusation threshold Z and the desired error rates ε_1 and ε_2 are also shown.

- For fixed m and Z, increasing c beyond c_0 affects only the distribution of S. The left curve remains unchanged and hence the FP rate is independent of the coalition size. This is compatible with the definition of ε_1-soundness (Definition 1).

- The FP error rate is determined by one parameter: $Z/\sqrt{m}$. Hence Z must be chosen as $Z \propto \sqrt{m}$ as far as the dependence on c_0 is concerned.

- When c increases, the rightmost curve becomes narrower and shifts to the left. In order to prevent the center of this curve ($\propto \sqrt{m}/c$) from crossing the threshold line, we need $m = \Omega(c^2)$. Together with the previous point, this illustrates the need for the proportionalities $m \propto c_0^2$, $Z \propto c_0$ in the Tardos scheme.

More precise results are derived next.

Lemma 3: A sufficient condition for ε_1-soundness and (c_0, ε_2)-completeness is given by

$$Z \in \left[\tilde{\sigma}_j \sqrt{m}\, G_1^{\mathrm{inv}}(\varepsilon_1), \quad \frac{\tilde{\mu}}{c_0} m + \frac{\tilde{\sigma}}{c_0} \sqrt{m}\, G_2^{\mathrm{inv}}(\varepsilon_2) \right]. \tag{63}$$

Proof of Lemma 3: The left boundary directly follows from the requirement $G_1(Z/\sigma_j) \leq \varepsilon_1$, using the notation (59). The right boundary follows from the requirement $G_2([c_0 Z - \mu]/\sigma) \leq \varepsilon_2$. □

Note that $G_2^{\mathrm{inv}}(\varepsilon_2)$ is negative on the interval $\varepsilon_2 < 1/2$ for symmetric G_2, and that it monotonically increases as a function of ε_2.

The Z-interval (63) exists only for sufficiently large m. One can think of a region in the (u, v)-plane, with $u = \sqrt{m}$ and $v = Z$, bounded on the lower side by a line $v \propto u$ and on the upper side by a quadratic function of u. The linear and quadratic curve meet each other at $m = m_{\min}$,

$$m_{\min} = \tilde{\mu}^{-2} c_0^2 \left[\tilde{\sigma}_j\, G_1^{\mathrm{inv}}(\varepsilon_1) - \frac{\tilde{\sigma}}{c_0}\, G_2^{\mathrm{inv}}(\varepsilon_2) \right]^2. \tag{64}$$

This represents the smallest possible code length for which the interval (63) exists, and hence, by Lemma 3, the smallest possible code length for which the code is properly collusion resistant.

5.5 'Extremal' colluder strategy

Eq. (64) allows us to find the 'worst case' or 'extremal' colluder strategy. We define this as the strategy that causes the highest possible value of $m_{\min}$. Even though the colluders do not necessarily use this strategy, the content owner has to take into account that they *might* and has to adjust m accordingly. We make the following observations:

- One way of increasing $m_{\min}$ would be to make $\tilde{\mu}$ as small as possible. However, in Section 5.3 it was shown that the choice of strategy has negligible effect on $\tilde{\mu}$.

- Another way of increasing $m_{\min}$ is to make $\tilde{\sigma}_j$ and $\tilde{\sigma}$ as large as possible. Here the choice of strategy has a big impact. Both $\tilde{\sigma}_j$ and $\tilde{\sigma}$ are maximally large if the coalition outputs a '1' whenever possible.

We see that the 'extremal' strategy is to set $\psi_i(x_i) = 1$ for $x_i \neq 0$. (The marking condition enforces $\psi_i(0) = 0$.) It looks as if the colluders are incriminating themselves in those columns where $p_i < 1/2$. However, for each symbol they equally incriminate a fraction p_i of all the other users. The strategy derives its effectiveness from the large number of users who get accused along with the colluders.

Substitution of the extremal strategy into (53) and (58) with $c = c_0$ gives

$$\tilde{\sigma}_j^2 = 1 - \frac{1}{\pi - 4t'} \int_t^{1-t} \mathrm{d}p \, p^{-1/2}(1-p)^{c_0-1/2} = 1 - \frac{1}{\sqrt{\pi c_0}} + \mathcal{O}(c_0^{-1})$$

$$\tilde{\sigma}^2 + \tilde{\mu}^2 = c_0 - \frac{c_0^2}{\pi - 4t'} \int_t^{1-t} \mathrm{d}p \, p^{1/2}(1-p)^{c_0-3/2} = c_0[1 - \frac{1}{2\sqrt{\pi c_0}} + \mathcal{O}(\frac{1}{c_0})] \qquad (65)$$

5.6 Final step in the proof of Theorem 2

If S_j and S have a Gaussian distribution, then the functions G_1, G_2 become error functions, and we have

$$G_1^{\mathrm{inv}}(\varepsilon_1) = \sqrt{2} \, \mathrm{Erfc}^{\mathrm{inv}}(2\varepsilon_1) \qquad ; \qquad G_2^{\mathrm{inv}}(\varepsilon_2) = -\sqrt{2} \, \mathrm{Erfc}^{\mathrm{inv}}(2\varepsilon_2). \qquad (66)$$

We obtain the following inequalities from (65),

$$\tilde{\sigma}_j < 1, \qquad ; \qquad \tilde{\sigma} < \sqrt{c_0}. \qquad (67)$$

These are independent of the choice of strategy function ψ_i. Likewise, from (57) we also obtain an inequality that is independent of ψ_i. This is done by taking only the negative part of the summand in (57) and setting $\psi_i = 1$. The result is

$$\tilde{\mu} > \frac{1 - 2\tau}{\pi}. \qquad (68)$$

We substitute (66) into (63) and (64), and then use the inequalities (67,68). This exercise shows that the choice (46) for m is larger than $m_{\min}$, as it should indeed be, and that the Z-interval (47) lies within the interval (63). This completes the proof of Theorem 2. $\qquad\square$

5.7 Proof of Corollary 3

We use the inequality [Wol]

$$\ln \frac{1}{x} \sqrt{\frac{2}{\pi}} > \left[\mathrm{Erfc}^{\mathrm{inv}} x \right]^2 \qquad (69)$$

to prove that the expression $\ln(1/\varepsilon_1\sqrt{2\pi})$ in (48) is larger than $[\mathrm{Erfc}^{\mathrm{inv}}(2\varepsilon_1)]^2$. Thus the following code length achieves Soundness and Completeness,

$$m = \frac{2\pi^2}{(1-2\tau)^2} c_0^2 \ln \frac{1}{\varepsilon_1\sqrt{2\pi}} \left[1 + \frac{1}{\sqrt{c_0}} \sqrt{\frac{\ln \varepsilon_2\sqrt{2\pi}}{\ln \varepsilon_1\sqrt{2\pi}}} \right]^2. \qquad (70)$$

We then neglect the term containing ε_2 with respect to 1, since it is of order $\mathcal{O}(1/\sqrt{c_0})$. This yields the value of m in Corollary 3. Finally, the value of Z in Corollary 3 follows by substituting this m into the left boundary in (47). $\qquad\square$

Remarks: In the regime $\varepsilon_1 \ll \varepsilon_2$, the fraction of logarithms in (70) is typically smaller than 0.1 (see Section 4). Hence, the asymptotic result is already approached for relatively small values of c_0.

In the case where ε_2 and ε_1 are coupled according to $\varepsilon_2 = \varepsilon_1^{c_0/4}$, as was done in Tardos' original construction, Corollary 3 does not hold, as we get $\ln \varepsilon_2 = (c_0/4) \ln \varepsilon_1$. Here the fraction of logarithms is not negligible and leads to a factor $(9/2)\pi^2$ in m instead of $2\pi^2$.

6 Summary

We have reevaluated the performance of the Tardos fingerprinting scheme by parameterizing its numerical constants and fixed functions. We have further modified the scheme by decoupling the desired false negative and false positive error probabilities. Using a proof technique similar to the one in [Tar2003], we have shown how short the code length can be with provable ε_1-soundness and (c_0, ε_2)-completeness. The main results of our study can be summarized as follows:

- Tardos' accusation function g_1^{T} is 'optimal' in the sense that it minimizes the provably sufficient code length for our particular choice of proof method.

- Tardos' probability distribution function $f(p)$ is 'optimal' in the same sense within a limited class of functions which has the form $p^{a-1}(1-p)^{b-1}$.

- For sufficiently large c_0 values, and ε_2 independent of ε_1, the code length can be reduced from Tardos' $100c_0^2 \ln \varepsilon_1^{-1}$ to approximately $4\pi^2 c_0^2 \ln \varepsilon_1^{-1}$.

- When $\varepsilon_2 \gg \varepsilon_1$, for instance for content distribution applications, our numerical results show that a code length $m < 46c_0^2 \ln \varepsilon_1^{-1}$ is achievable already for $c_0 > 9$.

- For sufficiently large $c_0 \gg 1$, the accusation sums S_j for the innocent user and S for the coalition have probability distributions which are very close to Gaussian—due to the Central Limit Theorem. If these distributions are perfectly Gaussian, then, in the case of independent ε_1, ε_2, a code length of $m \approx 2\pi^2 c_0^2 \ln \varepsilon_1^{-1}$ is sufficient for achieving ε_1-soundness and (c_0, ε_2)-completeness.

A Condition for Completeness

In this appendix we derive an upper bound on the expression $\mathbb{E}_{yXp}[e^{-\alpha_2 S}]$. The first part of the derivation is directly copied from [Tar2003], so we will not repeat it here. We start our analysis at the earliest point where the approach with general A, B, f and g deviates from [Tar2003].

From partial evaluation of the X-average (which involves the binomial distribution for each column of X separately) and from $|i : y_i = 1| \leq m$, it can be shown that

$$
\begin{aligned}
\mathbb{E}_{yXp}[e^{-\alpha_2 S}] &\leq \left[E_{0,0} + E_{1,c} + \sum_{x=1}^{c-1} \binom{c}{x} \max(E_{0,x}, E_{1,x}) \right]^m \qquad (71) \\
E_{0,x} &:= \mathbb{E}_p[p^x(1-p)^{c-x}] \\
E_{1,x} &:= \mathbb{E}_p\left[p^x(1-p)^{c-x} \exp\left(-\alpha_2\{xg_1(p) + [c-x]g_0(p)\} \right) \right] \\
&= \mathbb{E}_p\left[p^x(1-p)^{c-x} \exp\left(\alpha_2 g_1(p)\frac{cp-x}{1-p} \right) \right].
\end{aligned}
$$

The term $E_{1,c}$ is easily bounded,

$$
E_{1,c} = \mathbb{E}_p[p^c e^{-c\alpha_2 g_1}] \leq E_{0,c} - c\alpha_2 \mathbb{E}_p[p^c g_1] + c^2\alpha_2^2 \mathbb{E}_p[p^c g_1^2]. \qquad (72)
$$

Next it is proven that (for $x = 1 \ldots c - 1$) $\max(E_{1,x}, E_{0,x}) \leq E_{0,x}+$ some positive expression. To this end the inequality $e^u \leq 1 + u + u^2$ is again used, which holds for $u < 1.7$.

$$p^x(1-p)^{c-x} \exp\left(\alpha_2 g_1(p)\frac{cp-x}{1-p}\right) \leq p^x(1-p)^{c-x}\left\{1 + \alpha_2 g_1(p)\frac{cp-x}{1-p} + \alpha_2^2\left[g_1(p)\frac{cp-x}{1-p}\right]^2\right\}$$

$$+\Theta\left(\alpha_2 g_1(p)\frac{cp-x}{1-p} - 1.7\right)(1-p)^{c-x}\exp\left[\alpha_2 g_1(p)\frac{cp-x}{1-p}\right]. \tag{73}$$

Here the term with the step function Θ ensures that the right-hand side is always larger than the left-hand side, even if the expression in the exponent exceeds 1.7, which may happen for $p \in (x/c, 1-t)$ if α_2 is not very small. Taking the expectation value of (73), we get

$$E_{1,x} \leq E_{0,x} + \alpha_2 K_{1,x} + \alpha_2^2 K_{2,x} + R_x \tag{74}$$

with

$$K_{1,x} := \mathbb{E}_p\left[p^x(1-p)^{c-x}g_1\frac{cp-x}{1-p}\right]$$

$$= -pfg_1 \cdot p^x(1-p)^{c-x}\Big|_{p=t}^{1-t} + \int_t^{1-t} dp\, p^x(1-p)^{c-x}\frac{d}{dp}(pfg_1) \tag{75}$$

$$K_{2,x} := \mathbb{E}_p\left[p^x(1-p)^{c-x}\left\{g_1(p)\frac{cp-x}{1-p}\right\}^2\right] \geq 0 \tag{76}$$

$$R_x := \mathbb{E}_p\left[\Theta\left(\alpha_2 g_1(p)\frac{cp-x}{1-p} - 1.7\right)(1-p)^{c-x}\exp\left[\alpha_2 g_1(p)\frac{cp-x}{1-p}\right]\right] \geq 0. \tag{77}$$

Now we have to upper bound $K_{1,x}$ by a nonnegative expression. Here we depart from [Tar2003]. Tardos makes a very specific choice for the f and g_1 function, namely $pf(p)g_1(p) =$ constant. We keep the derivation as general as we can. For the moment we simply assume that we can find tight bounds $K_x^{\text{bound}} \geq 0$ such that

$$K_{1,x} \leq K_x^{\text{bound}}. \tag{78}$$

Then we have

$$\max(E_{0,x}, E_{1,x}) \leq E_{0,x} + \alpha_2 K_x^{\text{bound}} + \alpha_2^2 K_{2,x} + R_x. \tag{79}$$

Substituting (79) and (72) into (71) we get

$$\mathbb{E}_{yXp}[e^{-\alpha_2 S}] \leq \left[\sum_{x=0}^c \binom{c}{x}E_{0,x} - c\alpha_2\mathbb{E}_p[p^c g_1] + \alpha_2\sum_{x=1}^{c-1}\binom{c}{x}K_x^{\text{bound}} + \alpha_2^2\sum_{x=1}^c\binom{c}{x}K_{2,x}\right.$$

$$\left. + \sum_{x=1}^{c-1}\binom{c}{x}R_x\right]^m. \tag{80}$$

The $E_{0,x}$ term contains a sum over the binomial distribution and simply yields $\mathbb{E}_p[1] = 1$. The $K_{2,x}$ term is bounded as follows,

$$\sum_{x=1}^c \binom{c}{x}K_{2,x} \leq \mathbb{E}_p\left[\frac{g_1^2}{(1-p)^2}\sum_{x=0}^c\binom{c}{x}p^x(1-p)^{c-x}(x-cp)^2\right] = c\mathbb{E}_p[g_1^2\frac{p}{1-p}] = c\nu. \tag{81}$$

Next we bound the R_x term. Note that the step function in (77) for fixed p is nonzero only if $x \leq x_{\max}$, where

$$x_{\max} := \left\lfloor c(1-t) - 1.7\frac{1-t}{\alpha_2 g_1(t)}\right\rfloor < c. \tag{82}$$

Furthermore, we note that the function multiplying the step function in (77) is monotonically decreasing as a function of p, provided that α_2 is 'small enough'. (This statement is made more

accurate in Section 4). This means that the expectation value $\mathbb{E}_p[\cdot]$ can be upper bounded by evaluating the integrand at the point $p = p_x^*$, the smallest value of p where the step function is nonzero,

$$R_x \leq (1 - p_x^*)^{c-x} \exp\left[\alpha_2(cp_x^* - x)\frac{g_1(p_x^*)}{1 - p_x^*}\right] \qquad \text{with} \qquad \alpha_2(cp_x^* - x)\frac{g_1(p_x^*)}{1 - p_x^*} = 1.7. \quad (83)$$

We introduce a numerical constant β such that $g_1(p_x^*) \geq (1 - p_x^*)^\beta$. (In Tardos' case $\beta = 1/2$.) From the definition of p_x^* (83) it then follows that

$$(1 - p_x^*) \leq \left[\frac{\alpha_2}{1.7}(cp_x^* - x)\right]^{1/(1-\beta)} < \left[\frac{\alpha_2}{1.7}(c - x)\right]^{1/(1-\beta)}. \quad (84)$$

This gives us the following bound on R_x for $x < c$:

$$R_x < e^{1.7}\left[\frac{\alpha_2}{1.7}(c - x)\right]^{\frac{c-x}{1-\beta}}. \quad (85)$$

The R_x-sum in (80) can then be bounded as

$$\sum_{x=1}^{c-1}\binom{c}{x}R_x \quad \leq \quad \sum_{x=0}^{c-1}\binom{c}{x}R_x < e^{1.7}\sum_{x=0}^{x_{\max}}\binom{c}{x}\left[\frac{\alpha_2}{1.7}(c - x)\right]^{\frac{c-x}{1-\beta}} \leq e^{1.7}\sum_{x=0}^{x_{\max}}(\frac{ce}{c-x})^{c-x}\left[\frac{\alpha_2}{1.7}(c - x)\right]^{\frac{c-x}{1-\beta}}$$

$$= \quad e^{1.7}\sum_{x=c-x_{\max}}^{c}(\frac{ce}{x})^x\left[\frac{\alpha_2}{1.7}x\right]^{\frac{x}{1-\beta}} < e^{1.7}\sum_{x=c-x_{\max}}^{c}\left[e(\alpha_2 c/1.7)^{\frac{1}{1-\beta}}\right]^x$$

$$= \quad e^{1.7}\left[e(\alpha_2 c/1.7)^{\frac{1}{1-\beta}}\right]^{c-x_{\max}}\frac{1 - \left[e(\alpha_2 c/1.7)^{\frac{1}{1-\beta}}\right]^{x_{\max}+1}}{1 - \left[e(\alpha_2 c/1.7)^{\frac{1}{1-\beta}}\right]}$$

$$=: \quad e^{1.7}\Delta^{c-x_{\max}}\frac{1 - \Delta^{x_{\max}+1}}{1 - \Delta} < e^{1.7}\Delta^{c-x_{\max}}\frac{1}{1 - \Delta} \quad (86)$$

where we have introduced the abbreviation Δ for the small[6] value $e(\alpha_2 c/1.7)^{1/(1-\beta)}$. Summarizing, from (80) we obtain

$$\mathbb{E}_{yXp}[e^{-\alpha_2 S}] < \left[1 - c\alpha_2\mathbb{E}_p[p^c g_1] + \alpha_2\sum_{x=1}^{c-1}\binom{c}{x}K_x^{\text{bound}} + \nu c\alpha_2^2 + e^{1.7}\Delta^{c-x_{\max}}\frac{1}{1 - \delta}\right]^m. \quad (87)$$

Finally we impose the following condition on the parameters t, α_2:

$$1 - c\alpha_2\mathbb{E}_p[p^c g_1] + \alpha_2\sum_{x=1}^{c-1}\binom{c}{x}K_x^{\text{bound}} + \nu c\alpha_2^2 + e^{1.7}\Delta^{c-x_{\max}}\frac{1}{1 - \Delta} < 1 - \frac{\alpha_2}{L}, \quad (88)$$

where $L > 0$ is a numerical constant. The satisfiability of this condition depends on the choice of f, g_1 and L. Given that the condition is satisfied, we have the upper bound

$$\mathbb{E}_{yXp}[e^{-\alpha_2 S}] < [1 - \alpha_2/L]^m \leq \exp(-\alpha_2\frac{m}{L}). \quad (89)$$

For given f and g_1, we will be interested in the smallest value of L that can be achieved. Tardos chose t and α_2 such that $L = 4$.

---

[6] As long as β does not deviate too much from the Tardos case, Δ is of order > 1 in the small parameter α_2.

B   Numerical results

Table 1 shows the results of the numerical experiments described in Section 4. The values A, B and t parametrize the code length, accusation threshold and p-axis cutoff, respectively. We also note, for the sake of completeness, that the auxiliary variables take the following values for the parameter sets indicated in the table: $1.001 \leq L/\pi \leq 1.026$, $1.48 \leq \alpha_1/\alpha_1^{\mathrm{T}} \leq 1.58$ and $0.05 \leq \alpha_2/\alpha_2^{\mathrm{T}} \leq 0.30$.

Table 1: Numerical results								
$R \downarrow$	$\rightarrow c_0$	10	15	20	30	40	60	80
0.02	A	41.31	41.26	41.16	40.99	40.85	40.66	40.54
	B	12.86	12.85	12.83	12.80	12.78	12.75	12.73
	t/t^{T}	3.26	2.25	1.72	1.00	0.88	0.50	0.27
0.04	A	42.80	42.47	42.21	41.85	41.59	41.27	41.06
	B	13.08	13.03	13.00	12.94	12.90	12.85	12.82
	t/t^{T}	2.96	2.17	1.44	1.14	0.77	0.52	0.43
0.06	A	43.95	43.41	43.03	42.50	42.17	41.73	41.46
	B	13.26	13.18	13.12	13.04	12.99	12.92	12.88
	t/t^{T}	3.41	2.28	1.53	1.04	0.65	0.49	0.37
0.08	A	44.93	44.22	43.72	43.07	42.65	42.13	41.80
	B	13.41	13.30	13.22	13.13	13.06	12.98	12.93
	t/t^{T}	3.16	2.27	1.71	1.04	0.77	0.62	0.33
0.10	A	45.80	44.93	44.34	43.58	43.09	42.48	42.11
	B	13.54	13.41	13.32	13.20	13.13	13.04	12.98
	t/t^{T}	3.31	1.94	1.60	1.23	0.77	0.46	0.39

C   The Gaussian approximation

Under some reasonable assumptions, we can regard the accusation sums as Gaussian-distributed stochastic variables. Here, we outline our assumptions and show that the Central Limit Theorem (CLT) is applicable under these conditions. We first note the complete column symmetry and column independence of both the code generation process and the accusation method. Given this symmetry, we argue (without providing a proof) that the best colluder strategy for generating the colluded copy is also symmetric, i.e. their output y_i is independent of the column index i and independent of the X_C entries in the other columns ($\neq i$). Note that the 'extremal' colluder strategy of Section 5.5 also complies with this assumption. Given column symmetry and mutual independence of the accusation values, under the assumption of a symmetric colluder strategy, the accusation sums S, S_j are sums of i.i.d. variables, and the Central Limit Theorem (CLT) is applicable.

We show that the domain of applicability of the CLT is large enough to encompass a sufficient part of the tail of the S_j and S distributions, so that the approximations made in Section 5.6 are justified. The error probability $\varepsilon_1 \geq 10^{-15}$ represents at most an '8-sigma' event, i.e. we are interested in the region of 8 standard deviations σ_j around the average of S_j.

First we determine the probability distribution of each separate accusation $U(X_{ji}, p_i)$, for innocent j, given that $y_i = 1$. We define, for infinitesimal $\triangle u$,

$$\mathbb{P}[u \leq U \leq u + \triangle u] = \varphi(u)\triangle u. \tag{90}$$

We compute the conditional probability that $U = u$ given $X_{ji} = 1$. We write $\varphi(u|X = 1)\mathrm{d}u = f(p)\mathrm{d}p$, from which it follows that $\varphi(u|X = 1) = f(p)\mathrm{d}p/\mathrm{d}u$. Using $u = g_1(p)$, with g_1 defined in (1), we get $\varphi(u|X = 1) \propto 1/(1 + u^2)$. Applying the same reasoning to the case $X_{ji} = 0$, with $u = g_0(p)$, yields $\varphi(u|X = 0) \propto 1/(1 + u^2)$. From the conditional probability we obtain $\varphi(u)$ by

multiplying with the probability that the event $X = 1$ (or $X = 0$) occurs,

$$\varphi(u) = p\varphi(u|X=1) = (1-p)\varphi(u|X=0) \propto \frac{1}{(1+u^2)^2}. \tag{91}$$

Here we have used the fact that $p = 1/(1+u^2)$ for $X = 1$ and $p = u^2/(1+u^2)$ for $X = 0$. Thus the tail of the probability distribution has a $1/u^4$ power law behavior.

Next we argue that the number of contributing terms to S_j (almost m terms for the optimal colluder strategy) is sufficiently large for the CLT to cover 8 sigmas. For a distribution with vanishing third cumulant and with $\mathbb{E}[u^4] < \infty$, it is known (see e.g. [Baz2005]) that the region of convergence for the CLT, expressed in sigmas, is given by

$$\#\mathrm{sigmas} = \left(\frac{24\kappa_2^2}{\kappa_4}\right)^{1/4} N^{1/4}, \tag{92}$$

where N is the number of variables summed, and κ_j stands for the $j'th$ cumulant. Our distribution (91) satisfies the requirement $\mathbb{E}[u^4] < \infty$ because $\varphi(u)$ is defined on the finite interval $(-\sqrt{\frac{1-t}{t}}, -\sqrt{\frac{t}{1-t}}) \cup (\sqrt{\frac{t}{1-t}}, \sqrt{\frac{1-t}{t}})$. We have $\kappa_2 = 1 + \mathcal{O}(\sqrt{t})$ and $\kappa_4/\kappa_2^2 = 4/(\pi\sqrt{t}) + \mathcal{O}(1)$. Substitution into (92), with $N = m = 2\pi^2 c_0^2 \ln\varepsilon_1^{-1}$, gives

$$\#\mathrm{sigmas} = (12\pi^3 \ln\varepsilon_1^{-1})^{1/4} t^{1/8} \sqrt{c_0} \approx 5.2 c_0^{3/8}, \tag{93}$$

where we have used $\varepsilon_1 \approx 10^{-15}$ and $t^{\mathrm{T}} = 1/(300 c_0)$. Hence, the 8-sigma point of the tail is correctly approximated by a Gaussian already at $c_0 \geq 4$.

Acknowledgments

References

[Baz2005]   M.Z. Bazant, *Random walks and diffusion*, MIT Lecture notes, 2005, http://hdl.handle.net/1721.1/35916

[BS1998]   D. Boneh, J. Shaw, *Collusion-secure fingerprinting for digital data*, IEEE Transactions on Information Theory, 44(5), 1897–1905 (1998).

[CFNB2000] B. Chor, A. Fiat, N. Naor, and B. Pinkas, *Tracing Traitors*, IEEE Transactions on Information Theory, 46(3), 893–910, 2000.

[CNN2004]   http://www.cnn.com/2004/SHOWBIZ/01/23/oscar.arrest/index.html

[DCI2007]   Digital Cinema Initiatives, LLC, *Digital Cinema System Specification V1.1*, 2007, http://www.dcimovies.com/DCI_DCinema_System_Spec_v1_1.pdf

[HHI2006]   M. Hagiwara, G. Hanaoka, H. Imai, *A Short Random Fingerprinting Code Against a Small Number of Pirates*, in M. Fossorier et al. (Eds.): AAECC 2006, LNCS 3857, pp. 193–202 (2006)

[HvLLT1998] H.D.L. Hollmann, J.H. van Lint, J-P. Linnartz, L.M.G.M. Tolhuizen, *On codes with the identifiable parent property*, Journal of Combinatorial Theory, 82, 121–133, 1998.

[PSS2003]   C. Peikert, A. Shelat, A. Smith, *Lower bounds for collusion-secure fingerprinting*, in Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) 2003, pp. 472–479.

[RW2004]   R. Rob, J. Waldfogel, *Piracy on the High C's: Music Downloading, Sales Displacement, and Social Welfare in a Sample of College Students*, `http://ideas.repec.org/p/nbr/nberwo/10874.html`, 2004.

[SSW2001]   J.N. Staddon, D.R. Stinson, R. Wei, *Combinatorial properties of frameproof and traceability codes*, IEEE Transactions on Information Theory, 47(3), 1042–1049, 2001.

[Tar2003]   G. Tardos, *Optimal Probabilistic Fingerprint Coding*, in: Proceedings of the 35th Annual ACM Symposium on Theory of Computing, 2003, pp. 116– 125.

[Wol]   `http://functions.wolfram.com/GammaBetaErf/InverseErf/06/02/`.